

AMENDMENTS TO THE CLAIMS

This claims list supersedes all previous claims lists.

We claim:

1. (Currently Amended) A method for pipelining a table function in a database system, comprising:

a) performing a set up operation when the table function is called, the table function being a user-defined function that produces rows of data and can be used in selection, iteration, or aggregation database query language statements;

b) fetching a subset of output data from a data producer;

c) sending the subset of the output data to a first consumer of the output data, wherein the first consumer is the table function;

d) repeating steps b) and c) until all the output data has been fetched from the data producer.

2. (Original) The method of claim 1 in which the act of performing a setup operation comprises setting up a context object to maintain state.

3. (Original) The method of claim 1 in which the data producer comprises a second table function.

4. (Original) The method of claim 1 in which the subset of the output data comprises a single data object or row of data.

5. (Original) The method of claim 1 in which the subset of the output data comprises a plurality of data objects or rows of data.

6. (Original) The method of claim 1 further comprising:

e) performing a close operation after all the output data has been fetched from the data producer.

7. (Original) The method of claim 6 in which the act of performing the close operation comprises garbage collection operations.
8. (Original) The method of claim 7 in which the garbage collection operations comprises removal of a context object.
9. (Original) The method of claim 1 in which the table function executes in a different execution thread than the data producer.
10. (Original) The method of claim 1 in which the table function and the data producer execute from an identical execution thread.
11. (Original) The method of claim 1 in which a callback function is passed from the table function.
12. (Original) The method of claim 11 in which the callback function is executed on each subset of the output data fetched from the data producer.
13. (Original) The method of claim 1 in which the data producer comprises a dynamically configurable return type.
14. (Original) The method of claim 13 in which the dynamically configurable return type is established at compile time.
15. (Original) The method of claim 1 in which steps a) through d) are implemented within a database query language statement.
16. (Original) The method of claim 15 in which the database query language statement comprises SQL.
17. (Original) The method of claim 1 in which the subset of the output data is pipelined to a database query language statement.
18. (Original) The method of claim 17 in which a callback function is invoked for the subset of the output data.
19. (Original) The method of claim 18 in which the callback function filters inappropriate

data.

20. (Original) The method of claim 1 further comprising:

e) send the subset of the output data to a second consumer of the output data.

21. (Original) The method of claim 20 further comprising the step of determining whether the subset of the output data should be routed to the first consumer or the second consumer;

executing step c) if the subset of the output data should be routed to the first consumer;

and

executing step e) if the subset of the output data should be routed to the second consumer.

22. (Original) The method of claim 21 in which a partitioning definition is applied to determine whether the subset of the output data should be routed to the first consumer or the second consumer.

23. (Original) The method of claim 22 in which the partitioning definition comprises either hash or range based partitioning.

24. (Original) The method of claim 1 in which the first consumer processes the subset of the output data in parallel.

25. (Original) The method of claim 24 in which multiple slaves exist to process the subset of the output data.

26. (Original) The method of claim 25 further comprising the step of determining which of the multiple slaves operate upon the subset of the output data.

27. (Original) The method of claim 26 in which a partitioning definition is established to route the subset of the output data to an appropriate one of the multiple slaves.

28. (Original) The method of claim 27 in which the partitioning definitions comprises either hash or range based partitioning.

29. (Original) The method of claim 1 further comprising:
optimizing a query comprising the table function.
30. (Original) The method of claim 29 in which statistics for the table function are passed to an optimizer.
31. (Original) The method of claim 29 in which an optimizer self-determines statistics to optimize the query.
32. (Currently amended) A system for pipelining table functions, comprising:
- a) means for performing a set up operation when the table function is called, the table function being a user-defined function that can produce rows of data and can be used in selection, iteration, or aggregation database query language statements;
 - b) means for fetching a subset of output data from the table function;
 - c) means for sending the subset of the output data to a first consumer of the output data;
 - d) means for repeating steps b) and c) until all the output data has been fetched from the table function.
33. (Currently Amended) A computer program product comprising a computer useable medium having executable code to execute a process for pipelining table functions, the process comprising the steps of:
- a) performing a set up operation when the table function is called, the table function being a user-defined function that can produce rows of data and can be used in selection, iteration, or aggregation database query language statements;
 - b) fetching a subset of output data from the table function;
 - c) sending the subset of the output data to a first consumer of the output data;
 - d) repeating steps b) and c) until all the output data has been fetched from the table function.